Introduction to High Performance Computing

Alexander B. Pacheco

User Services Consultant
LSU HPC & LONI
sys-help@loni.org

HPC Training Fall 2013
Louisiana State University
Baton Rouge
September 4, 2013

1. What is HPC - Background and Defintions
2. Available HPC Resources
3. What is HPC@LSU?
4. HPC@LSU Services
5. Wrap-Up

- Gain understanding, mainly through the analysis of mathematical models implemented on computers.
- Construct mathematical models and quantitative analysis techniques, using computers to analyze and solve scientific problems.
- Typically, these models require large amount of floating-point calculations not possible on desktops and laptops.
- The field's growth drove the need for HPC and benefited from it.

- High Performance Computing (HPC) is computation at the forefront of modern technology, often done on a supercomputer.
- Acronym soup:
  1. HPC tasks are characterized as needing large amounts of computing power for short periods of time.
  2. High-throughput computing (HTC) tasks also require large amounts of computing, but for much longer times.
  3. High Productivity Computing Systems (HPCS) is a DARPA project for developing a new generation of economically viable high productivity computing systems.

- A supercomputer is a computer at the frontline of current processing capacity, particularly speed of calculation.
- A Supercomputer
  1. in the 70's used only a few processors
  2. in the 90's machines with thousands of processors appeared
  3. currently, massively parallel supercomputers with tens of thousands of "off-the-shelf" processors were the norm.
- Today, commodity PC's which you can purchase off-the-self have more than one core i.e. dual core, quad core processors.
- Smartphones and Tablets today have more processing power than 15 year old Supercomputers.
- What was a Supercomputer 15 years ago now sits on your desk, or even in your hand.

- Moore’s Law: number of transistors on integrated circuits doubles approximately every two years.
- So how is Supercomputing performance measured and who ranks them?

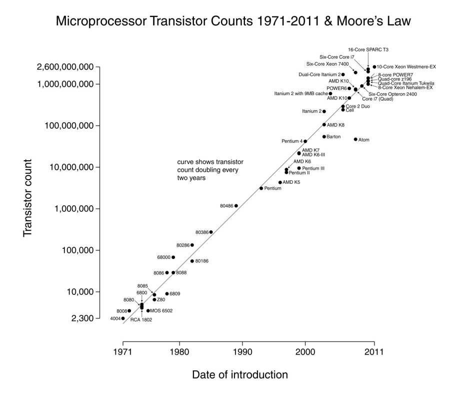

Source: `http://en.wikipedia.org/wiki/Moore's_law`

- The TOP 500 project ranks and details the 500 most powerful known computer systems in the world, published semi-annually
- Performance is measured in **FL**oating point **O**perations **P**er **S**econd (FLOPS or flop/s)
- The most powerful supercomputers nowdays
  - have more than a million cores
  - operate in petaflops (10^{15}) range.
- The fastest supercomputer as of June 2013 top500 list is Tianhe-2
  - Location: National University of Defense Technology
  - Nodes: 16000
  - Cores: 3120000
  - Peak Performance: 33.86 PFlop/s

$$GFLOPs = cores \times clock \times \frac{FLOPs}{cycle}$$

Most microprocessors today can do 4 FLOPs per clock cycle. Therefore a 2.5-GHz processor has a theoretical performance of 10 billion FLOPs = 10 GFLOPs.

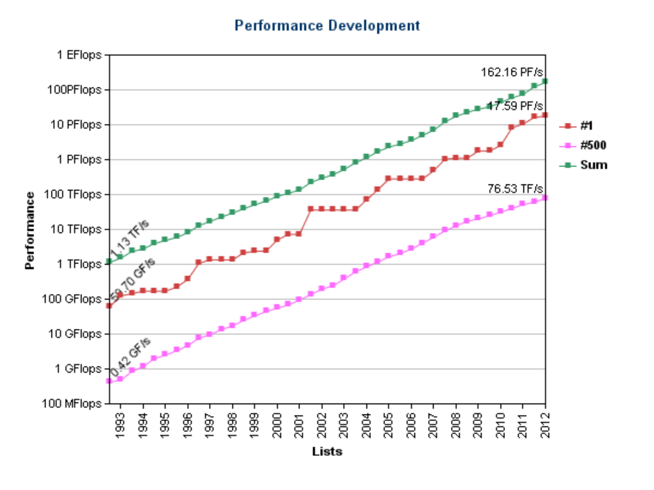

Source: `http://www.top500.org/statistics/perfdevel/`

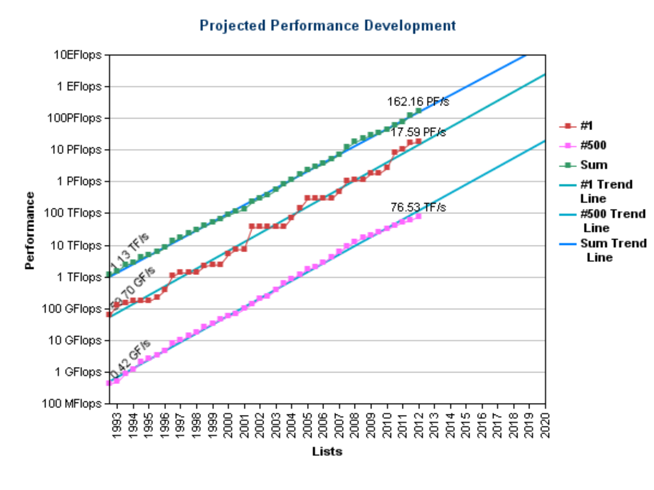

Source: `http://www.top500.org/statistics/perfdevel/`

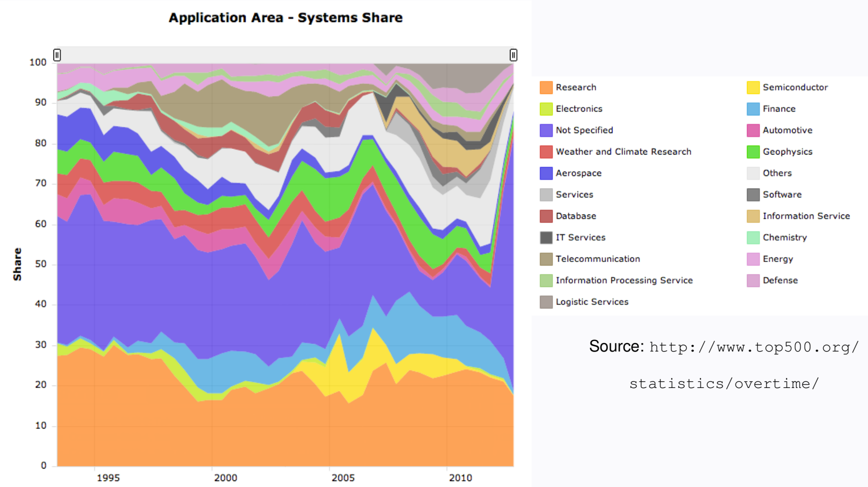

Source: `http://www.top500.org/statistics/overtime/`

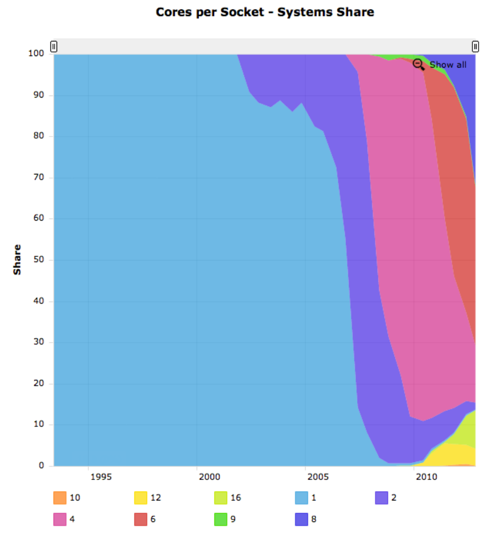

Source: `http://www.top500.org/statistics/overtime/`

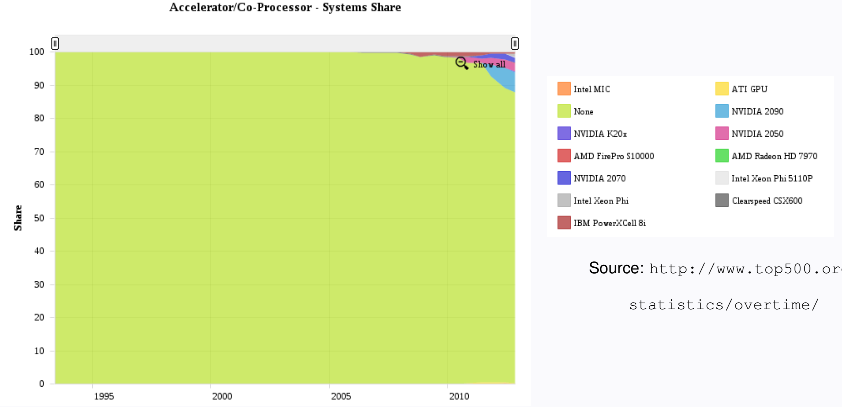

Source: `http://www.top500.org/statistics/overtime/`

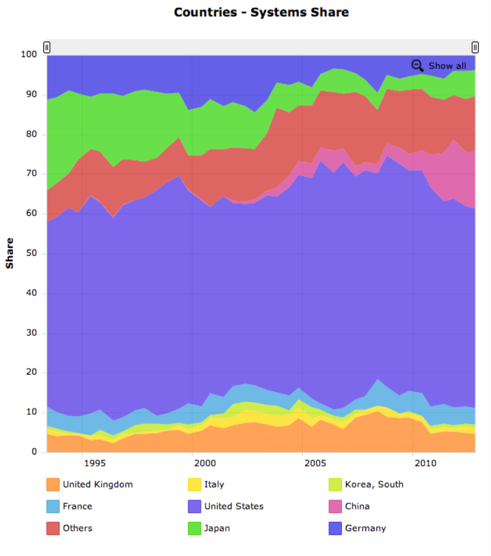

Source: `http://www.top500.org/statistics/overtime/`

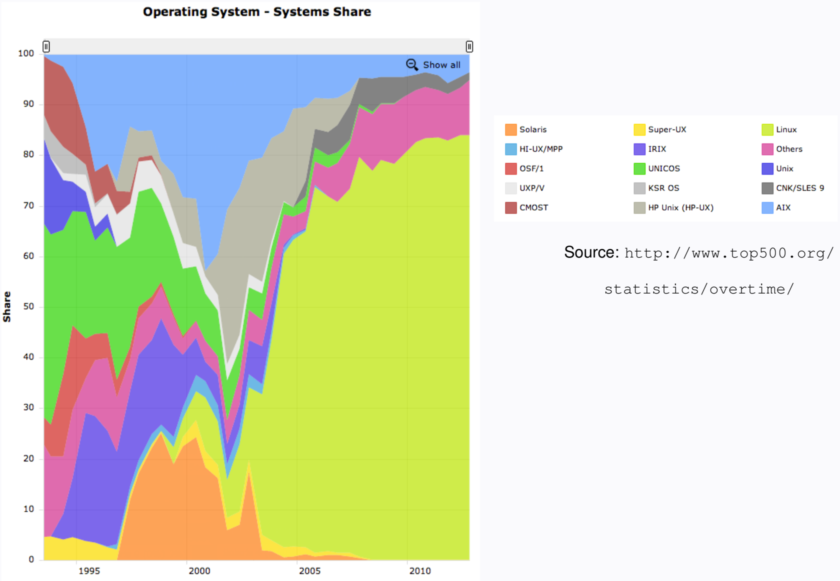

Source: `http://www.top500.org/statistics/overtime/`

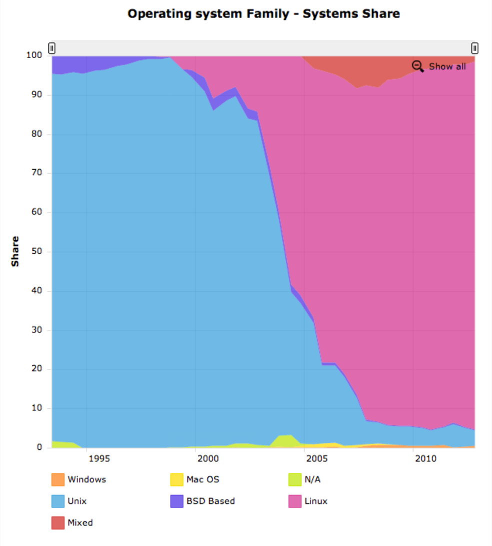

Source: `http://www.top500.org/statistics/overtime/`

Who uses these monstrous systems?

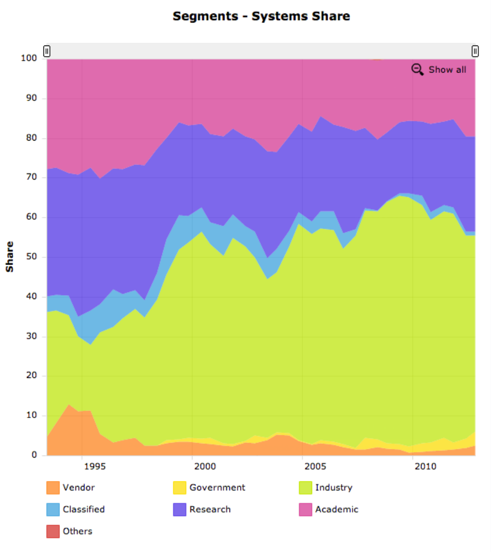

Source: `http://www.top500.org/statistics/overtime/`

- HPC may be the only way to achieve computational goals in a given amount of time
  - Size: Many problems that are interesting to scientists and engineers cannot fit on a PC usually because they need more than a few GB of RAM, or more than a few hundred GB of disk.
  - Speed: Many problems that are interesting to scientists and engineers would take a very long time to run on a PC: months or even years; but a problem that would take a month on a PC might only take a few hours on a supercomputer

- many calculations are carried out simultaneously
- based on principle that large problems can often be divided into smaller ones, which are then solved in parallel
- Parallel computers can be roughly classified according to the level at which the hardware supports parallelism.
  1. Multicore computing
  2. Symmetric multiprocessing
  3. Distributed computing
  4. Grid computing
  5. General-purpose computing on graphics processing units (GPGPU)

- using resources to solve a problem by dividing into many tasks, each of which is solved by one or more computers connected by a network.
- Condor
  - Work: run my application
  - Resources: a dedicated cluster of computers or idle computers on a university network
- Google
  - Work: process a query or "google search”
  - Resources: a lot of servers located worldwide
- SETI@HOME
  - Work: process signal data to find ET
  - Resources: a few servers and lots of PCs worldwide
- FOLDING@HOME
  - Work: simulates protein folding, computational drug design, and other types of molecular dynamics.
  - Resources: idle processing resources of thousands of PCs of volunteer who have installed the software on their systems.

- A subset of Distributed Computing
- Three point checklist by Ian Foster
  1. coordinates resources that are not subject to centralized control,
  2. using standard, open, general-purpose protocols and interfaces,
  3. to deliver nontrivial qualities of service.
- Condor across a university fails test 2
- Google has issues on tests 1 and 2
- SETI@HOME passes all three tests

I. Foster. *"What is the Grid? A Three Point Checklist"*, `http://dlib.cs.odu.edu/WhatIsTheGrid.pdf`

- Volunteer computing is a type of distributed computing in which computer owners donate their computing resources (such as processing power and storage) to one or more “projects”.
- **B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing (BOINC) is an example platform that supports over 40 projects (as of Dec 2012 $\sim$ 7.279 PFLOPS)
  1. SETI@Home (730 TFLOPS)
  2. MilkyWay@Home (1.6 PFLOPS)
  3. Einstein@Home (210 TFLOPS)
- Folding@Home is another example with 4.195 PFLOPS as of Dec 2009

Source: Wikipedia

What does HPC do?

- Simulation of Physical Phenomena
  - Storm Surge Prediction
  - Black Holes Colliding
  - Molecular Dynamics
- Data analysis and Mining
  - Bioinformatics
  - Signal Processing
  - Fraud detection
- Visualization
- Design
  - Supersonic ballute
  - Boeing 787 design
  - Drug Discovery
  - Oil Exploration and Production
  - Automotive Design
  - Art and Entertainment

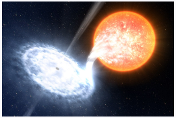

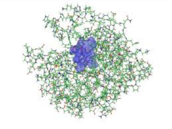

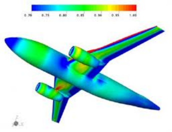

- Traditional Disciplines
  - Science: Physics, Chemistry, Biology, Material Science
  - Engineering
- Non Traditional Disciplines
  - Finance
    - Preditive Analytics
    - Trading
  - Humanities
    - Culturomics or cultural analytics: study human behavior and cultural trends through quantitative analysis of digitized texts, images and videos.

1. What is HPC - Background and Defintions
2. Available HPC Resources
3. What is HPC@LSU?
4. HPC@LSU Services
5. Wrap-Up

- National Level: **E**x**treme** **S**cience and **E**ngineering **D**iscovery **E**nvironment (xSEDE)
  - 5 year, $121M project supported by NSF
  - supports 16 supercomputers and high-end visualization and data analysis resources across the country.

XSEDE
Extreme Science and Engineering Discovery Environment

- State Level: **L**ouisiana **O**ptical **N**etwork **I**nitiative (LONI)
  - A state-of-the-art fiber optic network that runs throughout Louisiana and connects Louisiana and Mississippi research universities.
  - $40M Optical Network, 10Gb Ethernet over fiber optics.
  - $10M Supercomputers installed at 6 sites.

- University Level: LSU HPC resources available to LSU Faculty and their affiliates.
- ★ LONI and LSU HPC administered and supported by HPC@LSU

Outline

- Hardware Resources:
  - Currently manage 11 computing resources.
  - 7 LONI computing clusters
  - 4 LSU HPC computing clusters
- Available Software Stack
  - Communication Software
  - Programming support: Compilers and Libraries
  - Application Software
- User Services
  - Support: running jobs, software installation
  - Training:

Linux Clusters

Resource	Cluster	Peak TF/s	Location	Status	Login
LONI	QueenBee	50.7	ISB	Production	LONI
	Eric	4.7	LSU	Production	LONI
	Louie	4.7	Tulane	Production	LONI
	Oliver	4.7	ULL	Production	LONI
	Painter	4.7	LaTech	Production	LONI
	Poseidon	4.7	UNO	Production	LONI
	Satellite	4.7	Southern	Unknown	LONI
LSU HPC	Tezpur	15.3	LSU	Production	HPC
	Philip	3.5	LSU	Production	HPC
	SuperMike II	146 (CPU) 66 (GPU)	LSU	Production	HPC

AIX Clusters

Resource	Cluster	Peak TF/s	Location	Status	Login
LSU HPC	Pandora	6.8	LSU	Production	HPC

- Queen Bee
  - 668 nodes: 8 Intel Xeon cores @ 2.33 GHz
  - 8 GB RAM
  - 192 TB storage
- Other LONI clusters
  - 128 nodes: 4 Intel Xeons cores @ 2.33 GHz
  - 4 GB RAM
  - 9 TB storage
- Tezpur
  - 360 nodes, 4 Intel Xeon cores @ 2.33 GHz
  - 4 GB RAM
  - 32 TB storage
- Pandora
  - 8 Power7 nodes, 8 IBM Power7 processors @ 3.33 GHz
  - 128 GB RAM
  - 19 TB storage

- Philip
  - 37 nodes, 8 Intel Xeon cores @ 2.93 GHz
  - 24/48/96 GB RAM
  - 2 nodes, 12 Intel Xeon core @ 2.66GHz with hyperthreading with 3 Tesla 2070 GPU's each
  - Tesla M2070: 448 CUDA cores @ 1.15GHz and 5.25GB Total Memory
  - Shares storage with Tezpur

- Ranked 250 in Nov 2012 Top 500 List.
- 146 CPU TFlops and 66 double-precision GPU TFlops,
- 440 nodes, dual 8-core Intel Sandybridge Xeon cores @ 2.6 GHz
- 382 standard nodes with 32GB RAM (16 cores per node),
- 50 GPU nodes with 64GB RAM and dual NVIDIA Tesla M2090 6GB GPUs,
- 8 big memory nodes with 256GB RAM, capable of aggregation into a single virtual symmetric processing (vSMP) node using ScaleMP,
- Mellanox Infiniband QDR network of 2:1 over-subscription.

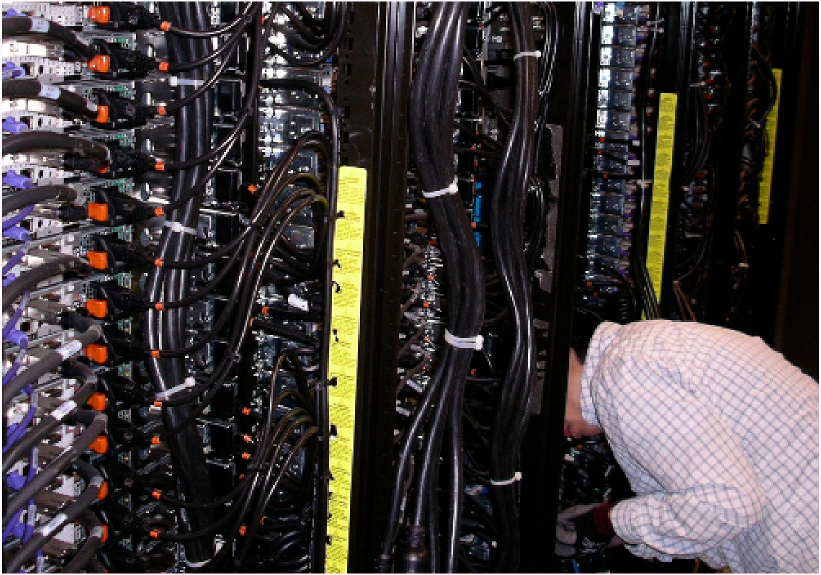


DataDirect
NETWORKS
S2A8500

LONI

- All faculty and research staff at a LONI Member Institution, as well as students pursuing sponsored research activities at these facilities, are eligible for a LONI account.
- Requests for accounts by research associates not affiliated with a LONI Member Institution will be handled on a case by case basis.
- For prospective LONI Users from a non-LONI Member Institution, you are required to have a faculty or research staff in one of LONI Member Institutions as your Collaborator to sponsor you a LONI account.

LSU HPC

- All faculty and research staff at Louisiana State University, as well as students pursuing sponsored research activities at LSU, are eligible for a LSU HPC account.
- For prospective LSU HPC Users from outside LSU, you are required to have a faculty or research staff at LSU as your Collaborator to sponsor you a LSU HPC account.

LONI Account

1. Visit `https://allocations.loni.org/login_request.php`.
2. Enter your **Institutional Email Address** and captcha code.
3. Check your email and click on the link provided (link is active for 24hrs only)
4. Fill the form provided
5. For LONI Contact/Collaborator field enter the name of your research advisor/supervisor who must be a Full Time Faculty member at a LONI member institution.
6. Click Submit button
7. Your account will be activated once we have verified your credentials.

LSU HPC Account

1. Visit `https://accounts.hpc.lsu.edu/login_request.php`.
2. Enter your **Institutional Email Address** and captcha code.
3. Check your email and click on the link provided (link is active for 24hrs only)
4. Fill the form provided
5. For HPC Contact/Collaborator field enter the name of your research advisor/supervisor who must be a Full Time Faculty member at LSU
6. Click Submit button
7. Your account will be activated once we have verified your credentials.

- LSU HPC and LONI systems are two distinct computational resources administered by HPC@LSU.

- Having an account on one does not grant the user access to the other.

What is an Allocation?

- An allocation is a block of computer time measured in core-hours (number of processing cores requested times the amount of wall-clock time used in hours).
- LONI users: All jobs need to be charged to valid allocation.
- LSU HPC users: Allocations are coming soon when SuperMike II is put into production.

Who can request an Allocation?

- Only Full Time Faculty member at LONI member institutions can act as Principle Investigators (PI) and request Allocations.
- Rule of Thumb: If you can sponsor user accounts, you can request allocations.
- Everyone else will need to join an existing allocation of a PI usually your advisor/supervision or course instructor (if your course requires a LONI account).

- Login to your LONI Profile at `https://allocations.loni.org`
- Click on "Request Allocation" in the right sidebar.
- Click "New Allocation" to request a New Allocation.
  1. Fill out the form provided.
  2. All requests require submission of a proposal justifying the use of the resources.
  3. Click "Submit Request" button.
- Click "Join Allocation" to join an existing Allocation.
  1. Search for PI using his/her email address, full name or LONI username
  2. Click "Join Projects" button associated with the PI's information.
  3. You will be presented with a list of allocations associated with the PI. Click "Join" for the allocation you wish to join.
  4. Your PI will receive an email requesting him to confirm adding you to the allocation.
  5. Please do not contact the helpdesk to do this.

Allocation Types

1. **Startup**: Allocations upto 50K SUs
   - Can be requested at any time during the year.
   - Reviewed and Approved by the LONI Resource Allocation Committee.
   - Only **two active** allocations per PI at any time.
   - Expired Allocations are considered active if the end date is in the future.
2. **Large**: Allocations between 50K - 4M SUs.
   - Reviewed and Approved by the LONI Resource Allocation Committee every Quater.
   - Users can have multiple Large Allocations.
   - A PI may have a total of 6M SUs active at any given time.

- LONI account

  `https://allocations.loni.org`
- LSU HPC account

  `https://accounts.hpc.lsu.edu`
- Newest cluster in production at LSU HPC is Pandora.
- Newest cluster at LSU HPC is SuperMike II is in user friendly mode.

- The default Login shell is bash
- Supported Shells: bash, tcsh, ksh, csh & sh
- Change Login Shell at the profile page

- LONI: Visit
  https://allocations.loni.org/user_reset.php
- LSU HPC: Visit
  https://accounts.hpc.lsu.edu/user_reset.php
- Enter the email address attached to your account and captcha code
- You will receive an email with link to reset your password, link must be used within 24 hours.
- Once you have entered your password, one of the HPC Admins need to approve the password reset.
- The Password approval can take anything from 10 mins to a few hours depending on the schedule of the Admins and also time of day
- You will receive a confirmation email stating that your password reset has been approved.

- Passwords should be changed as soon as your account is activated for added security.
- Password must be at least 12 and at most 32 characters long, must contain three of the four classes of characters:
  1. lowercase letters,
  2. uppercase letters,
  3. digits, and
  4. other special characters (punctuation, spaces, et cetera).
- Do not use a word or phrase from a dictionary,
- Do not use a word that can be obviously tied to the user which are less likely to be compromised.
- Changing the password on a regular basis also helps to maintain security.
- http://www.thegeekstuff.com/2008/06/the-ultimate-guide-for-creating-strong-passwords/
- http://en.wikipedia.org/wiki/Password_policy

- Shared Memory Programming - OpenMP
  - Good for programs that exhibit *data parallelism*.
  - Managed by compiler via special programming statements.
- Distributed Memory Programming - MPI
  - Good for programs that exhibit *task parallelism*.
  - Managed by programmer with library function calls.
- Hybrid Programming - OpenMP + MPI

- Compilers
  - Intel Fortran and C/C++
  - GNU compiler suite
  - Portland group Fortran and C/C++
  - CUDA
- Scripting languages
  - Perl, Python, BASH, TCSH, TCL/TK
- Numerical and utility libraries
  - FFTW, HDF5, NetCDF, PetSc, Intel MKL
- Debugging and Profiling Tools
  - DDT, TAU, TotalView

- Some things are common in scientific codes
  - Experts have developed and optimized methods for things like
    1. Matrix Operations
    2. Fast Fourier Transform
  - You do not need to reinvent the wheel, makes use of work done by the experts
- Many Scientific Libraries are available
  1. Linear Algebra: BLAS, ATLAS
  2. Linear Solvers: Scalapack, SuperLU, HYPRE, Intel MKL
  3. Fast Fourier Transform: FFTW, Intel MKL
  4. Boost

- Reading and Writing data is another common problem
  - ASCII: portable but slow
  - Binary: fast but not portable across machine architectures
  - what about metadata
- Common HPC I/O libraries
  - HDF: Hierarchial Data Format
  - NetCDF: Network Common Data Format
  - Manage format conversions between machines, can be annotated by metadata.
  - Used by many application, such as third-party visualization applications.

- Get your hands dirty:
  - Roll your own code
  - Install a source code application release
  - Modify an existing code
  - Understand what the code does for and to you
- ELSE, use an existing or installed packages and hope it satisfies all your research needs.

- Quantum Chemistry
  - Gaussian, GAMESS, NWCHEM, CPMD
- Molecular Dynamics
  - Amber, Gromacs, LAMMPS, NAMD
- Engineering
  - Fluent (LSU only)
- Mathematics and Statistics
  - Matlab (LSU only), Mathematica (LSU only), Octave, R
- Visualization
  - GaussView, VisIt, VMD, GNUPLOT

- Consulting Services
  - Usage Problems, Program Optimization, Software Installation, Software development advice.
- User Guides
  - HPC: `http://www.hpc.lsu.edu/docs/guides.php#hpc`
  - LONI: `http://www.hpc.lsu.edu/docs/guides.php#loni`
- Documentation: `https://docs.loni.org`
- Online Courses: `https://docs.loni.org/moodle`
- Contact us
  - Email ticket system: sys-help@loni.org
  - Telephone Help Desk: 225-578-0900
  - Instant Messenger: lsuhpchelp (AIM, Yahoo Messenger, Google Talk)

- Weekly Trainings
  1. Introductory: User Environment, Linux
  2. Programming: Shell Scripting, Perl, Python, MPI, OpenMP
  3. Software Development: Debugging, Profiling, Make, Subversion
  4. Software Applications: Molecular Dynamics, Computational Chemistry & Biology, Octave, MatLab
- Workshops
  1. Programming: Fortran, C, C++
  2. Parallel Programming: MPI, OpenMP, GPU, OpenACC
  3. Support Workshops organized through other Departments and Supercomputing centers.
- User Symposium.

♦ Held in June 2012 & 2013, Researchers from various LONI institutions presented their research via invited talks and poster sessions.

- HPC initiative is intended to stimulate and facilitate wider usage of HPC across the private sector to propel productivity, innovation and competitiveness.
- Goals
  1. Analyze the economic rationale for sustaining U.S. leadership in HPC, especially the impact upon manufacturing, services, business, and state-of-the-art research capabilities
  2. Identify key private sector HPC applications needs and priorities
  3. Identify workforce education and training needs to integrate HPC in the private sector
  4. Foster public-private sector partnerships to better leverage resources and expertise to help overcome barriers to more widespread private sector usage

---

http://www.compete.org/about-us/initiatives/hpc

The FUTURE is now.

Whatever happens in supercomputing now will be in your desktop in 10-15 years.

Having supercomputing experience now will keep you ahead of the curve when things get to the desktop in a decade or two.

The End

Any Questions?

Next Week:
Introduction to Linux

Survey:
`http://www.hpc.lsu.edu/survey`